

Error Control for Exponential Integration of the Master Equation

Katharina Kormann* Shev MacNamara†

October 12, 2016

Abstract

Error estimates for the numerical solution of the master equation are presented. Estimates are based on adjoint methods. We find that a good estimate can often be computed without spending computational effort on a dual problem. Estimates are applicable to both settings with time-independent, and time-dependent propensity functions. The Finite State Projection algorithm reduces the dimensionality of the problem and time propagation is based on an Arnoldi exponential integrator, which in the time-dependent setting is combined with a Magnus method. Local error estimates are devised for the truncation of both the Magnus expansion and the Krylov subspace in the Arnoldi algorithm. An issue with existing methods is that error estimates for truncation of the state space depend on measuring a loss of probability mass in a way that is not usually compatible with the approximation of the exponential. We suggest an alternative error estimate that is compatible with a Krylov approximation to the matrix exponential. Finally, we apply the new error estimates to develop an adaptive simulation algorithm. Numerical examples demonstrate the benefits of the approach.

Keywords: Master equation, Arnoldi method, adjoint method, chemical kinetics

AMS: 65F10, 65F60, 65L20, 65L70

1 Introduction

Markov processes offer a common mathematical framework with which to model applications as diverse as economics, chemical kinetics, single-molecule experiments, ecology, statistical mechanics, telecommunications, complex networks and more [34, 50, 2, 18, 26, 6, 19, 17, 11]. Estimating the probability of being

*Technische Universität München, Boltzmannstr. 3, D-85747 Garching, Germany (katharina.kormann@tum.de).

†The School of Mathematics and Statistics, University of New South Wales, UNSW, Sydney, Australia (s.macnamara@unsw.edu.au)

in a certain state is fundamental. Monte Carlo simulation of trajectories is one way to compute such estimates [14]. For example, in the context of chemical kinetics, the Gillespie Stochastic Simulation Algorithm and its variations are popular [21, 35].

Another approach is to directly solve a master equation (sometimes known as a forward Kolmogorov equation), which governs the evolution of the probability distribution associated with the discrete-state, continuous-time Markov process [54]. The master equation may be written as a linear system of ordinary differential equations

$$\frac{d}{dt}p = A(t)p, \quad p(0) = p_0. \quad (1)$$

Here $p \in \mathbb{R}^N$ is a probability vector: $p \geq 0$ component-wise, and $\|p\|_{\ell_1} = 1$. The i th component is the probability of state i . The off-diagonals of the propensity matrix $A \in \mathbb{R}^{N \times N}$ are nonnegative. The diagonals are such that each column of A sums to zero. The off-diagonals have the interpretation that $a_{ij}dt$ is the probability that flows from state j to state i in the next infinitesimal time interval dt [44, 51]. A simple example is the graph Laplacian on a line of nodes [53]. This article is concerned with the numerical solution of equation (1) when the matrix A has this special Markov structure.

When A is constant, finite and bounded, the solution to (1) is the matrix exponential,

$$p(t) = \exp(tA)p(0). \quad (2)$$

In some applications the number of states, N , and the size of the associated matrix, may be very large. This presents a computational challenge. Monte Carlo approaches are popular partly because of this challenge. Nevertheless, methods for the direct solution of the master equation have also been proposed, together with various approaches to the assessment of accuracy [16, 23, 15, 31, 25, 24, 43, 38, 37, 55].

One approach to this challenge is the Finite State Projection (FSP) algorithm [43], which is related to the method of finite sections [36]. A key idea is that, starting with an initial distribution of finite support, most of the probability stays localised nearby when propagating over a small time interval. Instead of working on the full state space, an FSP algorithm focuses computational effort on a subset of the full state space, which is smaller and therefore more efficient, but which still captures most of the probability, and so is still very accurate. With the assumption of exact matrix exponentials, and the law of conservation of probability, very attractive *a posteriori* error estimates have been derived [43]. However, usually an approximation to the matrix exponential in (2) is employed, such as an exponential integrator [41, 38, 37, 28, 47, 48, 32, 1]. These approximations do not exactly conserve probability so obtaining an estimate of the overall error when combining two or more such approximations is an issue. This issue serves as motivation for the present article, which shows how adjoint methods offer a mathematical framework that is more accommodating of combinations of approximations and that can be used for deriving overall error estimates.

When efficient estimates of the *a posteriori* error are available, adaptive step sizes can be used to compute the solution in an efficient way, and in a way that controls the global error. Adjoint methods find wide applications in engineering, and for ordinary differential equations *a posteriori* error estimates based on the adjoint method have been devised [13, 12, 10, 3]. In the context of the time-dependent Schrödinger equation, it has been shown that efficient error estimates can come from combining adjoint methods with conservation properties [33]. Here, we show how to adapt adjoint approaches such as appear in [13, 12, 3] to the master equation.

The outline of the article is as follows. In the next section, we devise *a posteriori* error estimates for the master equation. The estimates are based on a duality argument. Three different levels of approximation, with successively less dependence on the dual problem, are derived. These are general purpose estimates that may be combined with various numerical methods. Sec. 3 discusses truncation of the state space. We take special care of the fact that the matrix exponential cannot be considered to be computed exactly for numerical propagation. Next, we introduce a Magnus–Arnoldi propagator. Sec. 6 discusses how to compute a discretized error estimate for the chosen numerical integrator. Sec. 7 details one way to adaptively solve the master equation based on the *a posteriori* estimates previously introduced. We discuss the influence of the dual problem on the estimates, and compare different levels of approximation in Sec. 8. Sec. 9 presents numerical results, with examples from chemistry and immunology.

2 *A posteriori* error estimates

In this section, we devise *a posteriori* error estimates based on duality (cf. [13, 42]). These are generic estimates that may be combined with various numerical methods. Sec. 6 shows how to tailor these estimates to a specific numerical method based on Magnus–Arnoldi propagation and state space reduction.

The global error is estimated by a weighted sum of the local errors at each step. The estimate is of the form $\text{error} \leq (\text{condition number}) \times (\text{perturbation to data})$. This form highlights an analogy with the condition number, κ_A , of a matrix for a linear system, $Ax = b$, in which the relative error, ρ_x , in the solution x is bounded by $\rho_x \leq \kappa_A \rho_b$, where ρ_b is the relative error in (or perturbation to) the data b . In this analogy, the solution of a dual problem, which may be thought of as a sensitivity and which we denote by q below, plays a role similar to the condition number of a matrix, and the local numerical errors, denoted r below, play a role similar to perturbations to the data.

Consider the master equation (1). Let \tilde{p} denote a continuation of the discrete solution obtained from a numerical scheme that solves the perturbed equation

$$\frac{d}{dt}\tilde{p} = A(t)\tilde{p} + r(t), \quad \tilde{p}(0) = p_0 + R,$$

where $r(t)$ is the residual due to numerical error and R a perturbation in the

initial value. As in [13], we assume that the perturbations are bounded by a small constant. The error is

$$E = p(t_{\text{final}}) - \tilde{p}(t_{\text{final}})$$

at time t_{final} , where p is the exact solution to (1). We want to estimate some functional of the error. Let z be a normalized vector defining this functional. In order to estimate the functional $z^T E$, we consider the adjoint master equation

$$\frac{d}{dt}q = -A^T(t)q, \quad q(t_{\text{final}}) = z. \quad (3)$$

Note that the eigenvalues of the propensity matrix $A(t)$ all have non-positive real part. Therefore the forward problem (1) is stable. Then, also the adjoint equation is stable, as shown in [12, Thm. 4.1]. The fundamental theorem of calculus relates our error functional to the perturbation $r(t)$,

$$\begin{aligned} z^T E &= \int_0^{t_{\text{final}}} \left(\frac{d}{dt}q \right)^T E + q^T \left(\frac{d}{dt}E \right) dt + q(0)^T E(0) \\ &= \int_0^{t_{\text{final}}} (-A(t)^T q)^T E + q^T (A(t)p(t) - \tilde{p}(t) - r(t)) dt + q(0)^T R \\ &= - \int_0^{t_{\text{final}}} q(t)^T r(t) dt + q(0)^T R. \end{aligned} \quad (4)$$

Apply Hölder's inequality to estimate (4) by

$$|z^T E| \leq \int_0^{t_{\text{final}}} \|q(t)\|_{\infty} \|r(t)\|_{\ell_1} dt + \|q(0)\|_{\infty} \|R\|_{\ell_1}. \quad (5)$$

Many applications give rise to self-adjoint problems and ℓ_2 -norms are a natural choice, but the ℓ_1 -norm and ℓ_{∞} -norm are more suitable in our setting of the Master equation.

In order to approximate (4), consider the conservation properties of the adjoint equation (3), which can be deduced from the following properties of the master equation (cf. [54]):

$$\sum_k A_{kl} = 0 \quad \text{for each } l, \quad (6)$$

$$A_{kl} \geq 0 \quad \text{for } k \neq l. \quad (7)$$

The first property gives

$$\frac{d}{dt}q_k = - \sum_l (q_l - q_k) A_{lk}. \quad (8)$$

Let k_* be the index with $q_{k_*}(t) = \max_l q_l(t)$ at a certain point in time. The off-diagonals of A are nonnegative so from (8) we deduce $\frac{d}{dt}q_{k_*} > 0$. Hence, the ℓ_{∞} norm is decreasing when going backwards in time, i.e.,

$$\max_{t \in [0, t_{\text{final}}]} \|q(t)\|_{\infty} = \|q(t_{\text{final}})\|_{\infty} = \|z\|_{\infty}. \quad (9)$$

This further estimates $|z^T E|$ by

$$|z^T E| \leq \|z\|_\infty \left(\int_0^{t_{\text{final}}} \|r(t)\|_{\ell_1} dt + \|R\|_{\ell_1} \right). \quad (10)$$

An estimate without numerical evaluation of the dual problem At this point, we have derived three *a posteriori* error estimates: (4), (5), and (10). We make some observations about their relative merits. Usually, a disadvantage of adjoint methods is that they require the solution to the dual problem, which is expensive both to compute and to store in memory. For example, to apply estimate (4), we first need to solve the dual problem (3) for $q(t)$ and save this solution in memory for $(0, t_{\text{final}})$, so that we can then evaluate the integral in (4). In contrast, note that estimate (10) does *not* require the solution of a dual problem. This makes estimate (10) very attractive. Compared to estimate (4), estimate (10) is cheaper to compute but may overestimate the error. Whether or not it pays off to solve the dual is problem and parameter dependent. Below we show that (10) is particularly beneficial if we want to bound several functionals of the error at the same time.

Remark The result of equation (9) is based on the property $\sum_k A_{kl} = 0$ for each k . Truncating the state space, as described above, destroys this property. Then, estimates (5) and (10) hold all the same. Instead of (8), we have

$$\frac{d}{dt} q_k = - \sum_{l \in \mathcal{I}_1} (q_l - q_k) A_{lk} + \sum_{l \in \mathcal{I}_2} q_k A_{lk}.$$

Since $A_{lk} \geq 0$, $l \neq k$, and $q_k \geq 0$, we still have $\frac{d}{dt} q_{k^*} > 0$ for the index k^* where p takes its maximal value. This was the requirement for estimates (5) and (10).

A component-wise error bound Bounding the component-wise error is of interest [43]. For this purpose, consider functionals z defined by standard basis vectors $z = e_1, z = e_2, \dots$, where e_i is the i th column of the identity matrix, and observe that all have an ℓ^∞ norm of 1. Hence, by (10) the error in each component is bounded by $\int_0^{t_{\text{final}}} \|r\|_{\ell_1} dt + \|r(0)\|_{\ell_1}$. In this way, we have found a component-wise error bound that does not require knowledge of the adjoint equation. In that case, for a given tolerance $\varepsilon > 0$, if we control the error so that

$$\int_0^{t_{\text{final}}} \|r(t)\|_{\ell_1} dt + \|R\|_{\ell_1} \leq \varepsilon,$$

then the exact solution p is within ε of the numerical solution \tilde{p} :

$$\tilde{p} - \varepsilon e \leq p \leq \tilde{p} + \varepsilon e, \quad (11)$$

where $e = (1, \dots, 1)$. This estimate may be compared to estimate (2.9) in [43] or estimate (2.4) in [37], which is reproduced here in (12) for convenience. However,

those approaches focus on one primary source of error, rather than offering strategies to accommodate all sources of error. The approach presented here augments previous approaches by providing a unifying framework for including errors that arise from different sources. One common source of error comes from adaptively tracking the support, as we describe next.

3 Truncation of the state space

In many applications of master equations, the number of all possible states is enormous, whilst in practice, at any particular instant in time, only a small subset of these states contain most of the probability. This observation motivated numerical methods that adaptively track the support of the probability distribution, reducing memory requirements and computational time.

One approach to truncating the state space is as follows [43]. Let $A = A_T + A_R$ where A_T belongs to the truncated state space and A_R is the remainder. Hence, A_T consists of the square block indexed by the included states only, while A_R collects the rows and columns belonging to the dismissed states. Let \mathcal{I}_1 denote the set of states that are included, and let \mathcal{I}_2 denote the states that are not included. Then A_T is the submatrix of A with rows and columns indexed by the states of \mathcal{I}_1 . In MATLAB notation, $A_T = A(\mathcal{I}_1, \mathcal{I}_1)$. Instead of solving with the full matrix A , we solve

$$\frac{d}{dt}p^E(t) = A_T(t)p^E(t) = A(t)p^E(t) - A_R(t)p^E(t).$$

Assuming exact exponentials, it has been shown [43] that $\|p^E(t)\|_{\ell_1}$ decreases, and the amount by which it decreases from 1 is a measure of the component-wise error. Briefly, if $\|p(0)\|_{\ell_1} = 1$ and $\|p^E(t)\|_{\ell_1} = 1 - \delta$, then

$$p^E \leq p \leq p^E + \delta e, \quad (12)$$

where p is the exact solution.

An issue with this estimate is that we usually do not compute the matrix exponential exactly. Therefore, the ℓ_1 norm of the vector p will not be exactly conserved by the numerical propagator (see also the discussion in Sec. 10) so there is an issue with a state-space-truncation criterion of the form (12). To address this issue, observe that on a small time interval $[t, t + \Delta t]$, we can estimate the probability mass that exits the states in \mathcal{I}_1 by

$$\int_t^{t+\Delta t} A_R(\tau)p^E(\tau) d\tau. \quad (13)$$

Now, $p^E(t)$ is nonzero only for states included in \mathcal{I}_1 . Hence, computing $A_R p^E(t)$ is cheap because we only need to consider states in \mathcal{I}_2 that can be reached from \mathcal{I}_1 in one step. Sec. 9.2 compares this criterion (13) with (12), using an example in which the numerical method conserves the ℓ_1 norm. This shows good agreement between (13) with (12). Hence, (13) provides a viable alternative

criterion to estimate the error due to state space truncation when the estimation criterion (12) is unreliable because the numerical propagator does not exactly conserve the ℓ_1 norm.

4 A Krylov approximation to the matrix exponential

Solving the master equation (1) requires an approximation to the matrix exponential. In applications where the matrix is large and sparse (a common case in chemical master equations for example), the action of the exponential on a vector may be computed via a Krylov method, as we now describe [48].

Over one small time step Δt , consider the approximation of

$$p(t + \Delta t) = \exp(\Omega) p(t). \quad (14)$$

For master equations (1) with constant matrix A we have $\Omega = \Delta t \cdot A$ in mind. The case of time-varying $A(t)$ requires a more complicated choice of Ω , as we describe in the next section.

We now describe the Arnoldi method. For the purpose of this subsection we consider p to be the ‘correct’ function because we want to isolate the error solely due to the Arnoldi method. The Arnoldi process generates an orthogonal basis $V_s = [v_1 = p(t), v_2, \dots, v_s]$ of the s th order Krylov subspace $\mathcal{K}_s = \text{span}\{p(t), \Omega p(t), \dots, \Omega^{s-1} p(t)\}$ and an upper Hessenberg matrix H_s that represents the matrix $\Omega^{[m]}$ in the Krylov subspace,

$$\Omega V_s = V_s H_s + [H_{s+1}]_{s+1,s} v_{s+1}. \quad (15)$$

Here, $[H_{s+1}]_{s+1,s}$ denotes element $(s+1, s)$ of the projection matrix to the $(s+1)$ th Krylov subspace and v_{s+1} is the $(s+1)$ th vector of the basis V_s . We then approximate $p(t + \Delta t)$ by

$$p(t + \Delta t) \approx p^A(t + \Delta t) = V_s \cdot \exp(H_s) \cdot e_1 \cdot \|p(t)\|_{\ell_2}, \quad (16)$$

where $e_1 = (1, 0, \dots, 0) \in \mathbb{R}^s$.

Error control and stopping criteria for Krylov subspace methods have been studied by Saad [45] and Hochbruck & Lubich [28]. They consider the generalized residual, ρ_A , defined as

$$\rho_A = (\exp(\Omega))^{-1} p^A(t + \Delta t) - p(t). \quad (17)$$

From (15) they derive the following easy-to-compute expression for the residual:

$$\rho_A = \|p(t)\|_{\ell_2} \Delta t L_{s+1,s}^s [\exp(L^s)]_{s,1} \cdot v_{s+1}. \quad (18)$$

To see how the residual ρ_A fits into a *posteriori* estimate (4), we rearrange Eqn. (17),

$$p^A(t + \Delta t) = \exp(\Omega) (p(t) + \rho_A), \quad (19)$$

i.e. we solve for the exact right-hand-side, but with a perturbed initial value.

5 A Magnus integrator for a master equation with time-varying rates

When the matrix A is constant, the solution of (1) is the matrix exponential (2). When $A = A(t)$ is time-dependent, the evolution operator is more complicated. For sufficiently small intervals Δt ,

$$p(t + \Delta t) = \exp(\Omega(t)) p(t), \quad (20)$$

where the evolution operator is given by the Magnus expansion [39, 4], of which the first few terms are:

$$\Omega(t) = \int_t^{t+\Delta t} A(\tau) d\tau - \frac{1}{2} \int_t^{t+\Delta t} \left[\int_t^\tau A(\sigma) d\sigma, A(\tau) \right] d\tau + \dots$$

Here $[A, B] = AB - BA$ is the matrix commutator. All higher order terms in the expansion involve commutators. In the special case that the matrix commutes with itself, i.e. $[A(t_1), A(t_2)] = 0$, those commutators are all zero so the solution to (1) is $p(t + \Delta t) = \exp(\Omega(t))p(t)$ where $\Omega(t) = \int_t^{t+\Delta t} A(\tau) d\tau$. An example is when $A(t) = c(t)A_c$ for a scalar function $c(t)$ and a constant matrix A_c , in which case the integral in $\Omega(t) = A_c(\int_t^{t+\Delta t} c(\tau) d\tau)$ is scalar.

For a numerical integrator, the expansion is truncated to a certain order in Δt . Denoting the truncated Magnus matrix by $\Omega^{[m]}(t)$ with m being the order of the truncation, we propagate based on the formula,

$$p^M(t + \Delta t) = \exp(\Omega^{[m]}(t))p^M(t).$$

Naïve computation of successive terms in the Magnus expansion may be expensive. Fortunately, methods for the efficient computation of the terms in the truncated Magnus expansion up to a certain order have been devised by Iserles and Norsett [29]. Second-order and forth-order methods are most common in practice. Let θ_2 denote the second order truncation and θ_4 the additional term in a fourth order approximation. We use the optimized formulas presented by Blanes *et al.* [5]:

$$\theta_2 = \Delta t B^{(0)}, \quad \theta_4 = -(\Delta t)^2 [B^{(0)}, B^{(1)}],$$

where

$$B^{(j)}(t_n) = \frac{1}{(\Delta t)^{j+1}} \int_{-\Delta t/2}^{\Delta t/2} \tau^j A(t + \Delta t/2 + \tau) d\tau, \quad j = 0, 1, \dots$$

Error estimates for the truncated Magnus expansion Although finer theoretical estimates are known, it is usually more practical to assume that the first omitted term in the expansion dominates the error. If we propagate based on $\Omega^{[m]}$, the first order error term is

$$r_M(\tau) \approx \dot{\theta}_{m+2}(t)p^M(t), \quad t \leq \tau \leq t + \Delta t. \quad (21)$$

The residual r_M on interval $[t, t + \Delta t]$ enters the error estimate as

$$\int_t^{t+\Delta t} q(\tau)^T r_M(\tau) d\tau \approx \int_t^{t+\Delta t} q(\tau)^T \dot{\theta}_{m+2}(\tau) p^M(\tau) d\tau. \quad (22)$$

Since we do not know the values of $p(\tau)$ or $q(\tau)$ at any time but t and $t + \Delta t$ (or actually only for t when predicting Δt), we approximate the integral as

$$\begin{aligned} \int_t^{t+\Delta t} q(\tau)^T \dot{\theta}_{m+2}(\tau) p^M(\tau) d\tau &\approx q(t)^T \int_t^{t+\Delta t} \dot{\theta}_{m+2}(\tau) d\tau p^M(t) \\ &= q(t)^T \theta_{m+2}(t) p^M(t). \end{aligned} \quad (23)$$

This approximation is sufficiently accurate for our purpose because Δt is small and r_M is only approximated to first order. Finally, we can apply Holder's inequality to the approximated integral for estimates of type (5) or (10).

Remark Recall the two properties (6) and (7) of a matrix $A(t)$ for the master equation. The commutator $[A(t_1), A(t_2)]$ shares property (6) with $A(t)$ that columns sum to zero, but unlike $A(t)$, the commutator can have negative off-diagonal entries.

A special structure of master equations In general, the main costs of evaluating approximations to the Magnus expansion, such as the formulas for θ_4 , come from the evaluation of the commutators. However, in the context of master equations there is often a special structure to the time dependence, and this structure can help. The propensity matrix often has the form

$$A(t) = A_c + \sum_{l=1}^r f_l(t) A_l \quad (24)$$

where A_c and A_l are constant in time and the $f_l : \mathbb{R} \rightarrow \mathbb{R}$ are scalar functions. This special case arises when the propensity is of the *separable* form $\alpha(x, t) = \alpha(x)\alpha(t)$. The rate is separable into a product of two terms, one that depends only on the state, and the other that depends only on time. For example if the bimolecular reaction $A + B \xrightarrow{c(t)} C$ has time varying rate $c(t)AB$ then the propensity is of this form.

In this common case (24) we have

$$\begin{aligned} \theta_2 &= (\Delta t) \cdot A_c + \sum_{l=1}^r (g_l \cdot A_l), \\ \theta_4 &= (\Delta t)^2 \left(\sum_{l=1}^r (h_l \cdot [A_l, A_c]) + \sum_{l_1=1}^{r-1} \sum_{l_2=l_1+1}^r (g_{l_2} h_{l_1} - g_{l_1} h_{l_2}) [A_{l_1}, A_{l_2}] \right). \end{aligned}$$

where the scalar integrals are

$$g_l = \frac{1}{\Delta t} \int_0^{\Delta t} f_l(t + \tau) d\tau, \quad h_l = \int_{-\Delta t/2}^{\Delta t/2} \tau^j f_l(t + \Delta t/2 + \tau) d\tau, \quad l = 1, \dots, r. \quad (25)$$

Thus, the constant commutators $[A_c, A_l]$ and $[A_{l_1}, A_{l_2}]$ can be precomputed once and for all at the beginning. Then, in each time step, only *scalar* integrals are required.

Magnus integrators are accurate for master equations, even for large time steps For (20) to be valid, Moan and Niesen have shown that the following inequality is a sufficient (but not necessary) condition [40],

$$\int_t^{t+\Delta t} \|A(\tau)\|_{\ell_2} d\tau \leq \pi. \quad (26)$$

It is natural to wonder if the step size of a numerical method based on the Magnus expansion must be small enough to satisfy (26). (That would place an unacceptably severe restriction on the step size whenever the matrix has a large norm, which is often the case for large sparse matrices arising in master equations.) In practice, we find that we can solve the master equation with much larger time steps whilst maintaining stability (Figure 1). This good experience is known for the time-dependent Schrödinger equation and Hochbruck and Lubich [27] provide an explanation. Unfortunately, the same reasoning does not explain this good experience in the context of the master equation. Their reasoning is based on the fact that the Hamiltonian of the Schrödinger equation is self-adjoint and thus has only imaginary eigenvalues and can be transformed into a diagonal matrix with an orthogonal transformation. For the master equation, on the other hand, we usually do not have such a transformation.

6 An estimate for state space truncated Magnus–Arnoldi propagation

The numerical simulation with the Magnus–Arnoldi propagator based on a truncated state space has three sources of errors: the truncation of the state space, the truncation of the Magnus expansion, and the projection to the Krylov subspace. We view the discretization process in three steps and use the following notation: First we truncate the state space and denote the solution of the corresponding perturbed master equation by p^E . As a next step, we introduce a temporal grid, $0 = t_0, t_1, \dots, t_{N_t} = t_{\text{final}}$ and apply the truncated Magnus expansion for propagation on each interval. The solution of the corresponding perturbed master equation is denoted by p^{EM} . Finally, we use the Arnoldi algorithm to compute the matrix exponential and get the fully perturbed master equation with solution p^{EMA} .

As usual in numerical methods for ODEs, we only have estimates for the local error over a small time step, so we “glue” these local estimates together over many time steps, to arrive at a global estimate at the end of the full time interval. First we derive an estimate for a single time step. The restarted perturbed problems then have an initial value that is additionally perturbed due to errors from previous time steps.

At each time step $n \in \{1, \dots, N_t - 1\}$, we have the perturbed problem

$$\begin{aligned} \dot{p}^{EMA} &= Ap^{EMA} - A_R p^{EMA} + r_M(t), \quad t \in (t_n, t_{n+1}) \\ p^{EMA}(t_n) &= p(t_n) + \rho_A(t_n) + e(t_n). \end{aligned} \quad (27)$$

Here, $e(t_n) = p^{EMA}(t_n) - p(t_n)$ is the error accumulated from previous time steps and $-A_R p^{EMA}$, r_M , and ρ_A are perturbations due to state-space truncation, Magnus expansion truncation and the Arnoldi method, respectively. Perturbations due to the Magnus expansion, r_M , can be estimated by (21) and the perturbations due to the Arnoldi algorithm by (18). Note these residuals are computed with the truncated state space matrix A_T .

Applying the *a posteriori* error estimation theory from Sec. 2 to the short time interval, we get

$$q(t_{n+1})^T e(t_{n+1}) = \int_{t_n}^{t_{n+1}} q(t)^T (r_M(t) - A_R p^{EMA}(t)) dt + q(t_n)^T (\rho_A(t_n) + e(t_n)). \quad (28)$$

We picked a special functional of the error that is based on the dual problem at time t_{n+1} . In this way, the estimate from the previous step is the last term in the error estimate for the present step, $q(t_n)^T e(t_n)$, that covers errors from previous steps. So we get a chain of estimates,

$$\begin{aligned} q(t_1)^T e(t_1) &= \int_{t_0}^{t_1} q(t)^T (r_M(t) - A_R p^{EMA}(t)) dt + q(t_0)^T (\rho_A(t_0)), \\ q(t_{n+1})^T e(t_{n+1}) &= \int_{t_n}^{t_{n+1}} q(t)^T (r_M(t) - A_R p^{EMA}(t)) dt + q(t_n)^T (\rho_A(t_n) + e(t_n)), \\ n &= 1, \dots, N_t - 2, \\ z^T e(t_{N_t}) &= \int_{t_{N_t-1}}^{t_{N_t}} q(t)^T (r_M(t) - A_R p^{EMA}(t)) dt \\ &\quad + q(t_{N_t-1})^T (\rho_A(t_{N_t-1}) + e(t_{N_t-1})). \end{aligned} \quad (29)$$

The local error in each step is

$$\int_{t_n}^{t_{n+1}} q(t)^T (r_M(t) - A_R p^{EMA}(t)) dt + q(t_n)^T \rho_A(t_n). \quad (30)$$

Corresponding to (4), (5), and (10), we have the following three estimates:

$$z^T e(t) \leq \sum_{n=0}^{N_t-1} \left(\int_{t_n}^{t_{n+1}} q(t)^T (r_M(t) - A_R p^{EMA}(t)) dt + q(t_n)^T \rho_A(t_n) \right), \quad (31)$$

$$\begin{aligned} |z^T e(t)| &\leq \sum_{n=0}^{N_t-1} \left(\int_{t_n}^{t_{n+1}} \|q(t)\|_\infty \|r_M(t)\|_{\ell_1} dt \right. \\ &\quad \left. + \int_{t_n}^{t_{n+1}} \|q(t)\|_\infty \|A_R p^{EMA}(t)\|_{\ell_1} dt + \|q(t_n)\|_\infty \|\rho_A(t_n)\|_{\ell_1} \right), \end{aligned} \quad (32)$$

and

$$|z^T e(t)| \leq \|z\|_{\ell_\infty} \sum_{n=0}^{N_t-1} \left(\int_{t_n}^{t_{n+1}} \|r_M(t)\|_{\ell_1} dt + \int_{t_n}^{t_{n+1}} \|A_{RP}^{EMA}(t)\|_{\ell_1} dt + \|\rho_A(t_n)\|_{\ell_1} \right). \quad (33)$$

7 An adaptive simulation algorithm

In the previous section, we derived an *a posteriori* error estimate for propagation of the master equation. We now explain how this estimate can be used to control the error by *adaptively* choosing: the size of the state space, the size of the time step, and the dimension of the Krylov space. We have other parameters to control the state space truncation error and the Arnoldi error. Therefore, we primarily choose the step size to control the Magnus error. Then, we choose the state space as large as needed for the chosen step size, and choose the Krylov subspace large enough to keep the Arnoldi error small. This strategy works well when the time-dependent nature of the problem requires small time steps. In general, Krylov methods usually only give good approximations for a moderate dimension (say, 40) of the Krylov space. Therefore, we also aim to keep the time step small enough to keep the size of the Krylov space moderate. In a first approximation, ρ_A is proportional to Δt , so we may assume the ratio between Δt and the size s of the Krylov subspace is constant. Given s_n and Δt_n for step n , we have the following upper bound for the size of Δt_{n+1} ,

$$\Delta t_{n+1, \max} = \frac{s_{\max}}{s_n} \Delta t_n, \quad (34)$$

where s_{\max} is the maximum acceptable dimension of the Krylov subspace. Given a tolerance tol , the simulation based on estimate (33) is performed as:

1. *Try to reduce state space:* Every C steps (for some constant C) check if the state space can be reduced.
2. *Check if the state space needs to be increased:* Check whether or not $A_R(t_n + \frac{\Delta t_n}{2})(p(t_n) + p(t_{n+1}))/2 < \frac{\text{tol}}{t_{\text{final}}}$. If not, increase the state space.
3. *Propagate solution with adaptive Arnoldi method:* Propagate n th step. Choose size of Krylov subspace such that $\rho_A(t_n) < \frac{\text{tol} \cdot \Delta t_n}{t_{\text{final}}}$.
4. *Compute new step size according to Magnus residual:* Estimate the error due to the truncation of the Magnus expansion via (21) and compute the optimal time step for the next step according to the formula,

$$\Delta t_{n+1}^M = \left(\Delta t_n \cdot \frac{\text{tol}/t_{\text{final}}}{\|\theta_{m+2}(t_{n-1} + \Delta t_n) \cdot p(t_n)\|_{\ell_1}} \right)^{1/m} \cdot \Delta t_n. \quad (35)$$

5. *Keep the step size small enough for the Arnoldi method:* Check whether or not the step size Δt_{n+1}^M is acceptable and set

$$\Delta t_{n+1} = \min \left(\Delta t_{n+1}^M, \frac{s_{\max}}{s_n} \Delta t_n \right). \quad (36)$$

If we want to use estimate (31) or (32) instead of (33), then we additionally have to include the dual or its ℓ_∞ norm, respectively, in the error checking. Note that Step 4 is redundant if the matrix is constant, in which case we set $\Delta t_{n+1} = \frac{s_{\max}}{s_n} \Delta t_n$.

8 Dual problem

Estimates (31) and (32) require the solution of a dual problem while estimate (33) only requires the primal solution. The more information we use from the dual problem, the more accurate our estimate of the error will be. Thus larger — and fewer — time steps can be used. On the other hand, solving the dual problem is a computational cost. Hence, there is a trade-off between costs and benefits of computing the dual. Such trade-offs are familiar in the literature on adjoint methods, arising much more generally than merely in our applications.

Memory is an additional difficulty. For the estimates (31) and (32) we need primal and dual solutions at the same time. Therefore, we have to first solve the dual problem backwards in time, and then save this dual solution in memory, to then be ready to compute the solution of the forward problem. Note estimate (32) requires only the scalar $\|q(t)\|_\infty$ (and not the vector $q(t)$), which will usually fit in memory. Estimate (31) needs the full solution so it requires much more memory or checkpointing with additional requirements on computation.

To be competitive with estimate (33), we cannot spend too much effort on the dual problem so we only compute a rough approximation. However, we will, of course, not know a priori how the error in the dual behaves. In principle, we can appeal to the same theory as for the primal. Since we do not have the solution of the ‘dual of the dual’ — i.e., the primal — available, we may resort to using an estimate corresponding to (33) for the dual. Denoting the error in the dual solution at time zero by e_q and the residual by r_q , we find the following estimate for a functional defined by z ,

$$\begin{aligned} |z^T e_q| &\leq \int_{t_{\text{final}}}^0 p(t)^T r_q(t) dt + p(t_{\text{final}})^T q(t_{\text{final}}) \\ &\leq \int_{t_{\text{final}}}^0 \|p(t)\|_{\ell_1} \|r_q(t)\|_{\ell_\infty} dt + \|p(t_{\text{final}})\|_{\ell_1} \|r_q(t_{\text{final}})\|_{\ell_\infty} \\ &\leq \int_{t_{\text{final}}}^0 \|r_q(t)\|_{\ell_\infty} dt + \|r_q(t_{\text{final}})\|_{\ell_\infty}. \end{aligned} \quad (37)$$

So if we make sure that

$$\max_{t_n \leq t \leq t_{n+1}} \|r_q(t)\|_\infty \leq \frac{\varepsilon}{t_{\text{final}}} \Delta t \quad (38)$$

in each time step, then we bound the error at any time by ε . Of course, we should choose ε to be much larger than the tolerance to which we want to compute the solution. Choosing ε is therefore a familiar trade-off between the advantage of the more accurate error estimates that come from more accurate dual solutions, and the disadvantage of the cost of solving the dual.

Altogether, we conclude that estimate (33) is the most reliable estimate, because it does not depend on uncontrolled influences from the approximation of the dual. We invest extra effort solving the primal problem more accurately (in the subspace perpendicular to z) instead of solving an additional dual problem. This is the strategy of choice when z is unknown and also often when we are interested in several functionals at the same time.

9 Numerical experiments

We now present numerical results to complement our theoretical findings. Given a prescribed tolerance for the error in a given component, we simulate the solution of the master equation with our state-space truncated Magnus–Arnoldi propagator with adaptivity based on the three estimates. We compare effectiveness and efficiency of the three estimates. By effectiveness, we mean that the error is indeed below the tolerance. The closer the actual error is to the tolerance, the more efficient the estimate.

9.1 A test case with analytic solution

Consider a master equation, with time varying propensities,

$$\begin{aligned} \frac{d}{dt}p &= A(t)p = \left[\begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix} + \sin(t) \begin{pmatrix} -1 & -1 \\ 1 & 1 \end{pmatrix} \right] p, \\ p(0) &= (\sigma, 1 - \sigma)^T, \end{aligned} \quad (39)$$

for $\sigma \in [0, 1]$. The solution is

$$p(t) = \begin{pmatrix} \frac{1}{2} + \frac{1}{5} \cos(t) - \frac{2}{5} \sin(t) + \left(\sigma - \frac{7}{10}\right) \exp(-2t) \\ \frac{1}{2} - \frac{1}{5} \cos(t) + \frac{2}{5} \sin(t) - \left(\sigma - \frac{7}{10}\right) \exp(-2t) \end{pmatrix}. \quad (40)$$

Table 1 shows results for simulations of equation (39) with $\sigma = 1$ over a time interval of $t \in [0, 10]$. The experiments show that it does not pay off to compute the dual solution for this example. Table 1 shows all estimates are effective. It also shows the estimates are more efficient as more information from the dual is used, as one would expect from theory. Note that in this example, the only source of error is due to truncation in the Magnus expansion.

9.2 Chemical master equation

Reversible isomerization involves the conversion of a molecule between two different forms. Consider a chemical master equation (CME) describing this pro-

Table 1: Results for the simple master equation in (39) with the error checked by comparison of the numerical solution with the exact, analytic solution. Parameters: $t_{\text{final}} = 10$ and tolerance $= 10^{-3}$.

estimate	ℓ_∞ error	# time steps (+dual steps)
(31)	9.1×10^{-4}	62+100
(32)	3.6×10^{-4}	112+100
(33)	3.3×10^{-4}	131

cess:

$$X \xrightarrow{c_{12}} Y, \quad Y \xrightarrow{c_{21}} X. \quad (41)$$

We choose $c_{1,2}(t) = 1 + \sin(t)$ and $c_{2,1}(t) = 1 - \sin(t)$ and let the initial distribution be a binomial distribution $\mathcal{B}(x_k; N, p_0)$ with $N = 2000$ molecules and parameter $p_0 = (1/3, 2/3)^T$. Observing the conservation law, there are 2001 possible states $x_k = (k, N - k)$ for $k = 0, \dots, 2000$. The chemical master equation for (41) is,

$$\frac{d}{dt} P(t) = (A_0 + \sin(t)A_1) P(t), \quad (42)$$

where the k th element of P gives the probability that the system is in state x_k and

$$A_0^{(i,j)} = \begin{cases} -N & \text{if } i = j \\ j & \text{if } i = j - 1 \\ N - j & \text{if } i = j + 1 \\ 0 & \text{otherwise} \end{cases}, \text{ and } A_1^{(i,j)} = \begin{cases} N - 2j & \text{if } i = j \\ j & \text{if } i = j - 1 \\ -N + j & \text{if } i = j + 1 \\ 0 & \text{otherwise} \end{cases}. \quad (43)$$

The exact analytic solution to (42) is known (see Proposition 1 in [30]) to be

$$P_k(t) = \mathcal{B}(x_k; N, p(t)), \quad k = 0, \dots, N, \quad (44)$$

where the 1×2 parameter vector $p(t)$ is the solution of the reaction rate equation

$$\dot{p}(t) = A(t)p(t), \quad p(0) = p_0, \quad (45)$$

in which $A(t)$ is as in (39) of the previous example with $\sigma = 1/3$ and thus expression (40) gives the solution $p(t)$.

Table 2 compares the control strategies based on estimates (31), (32), and (33). As the error functional, we use the error in component 1150, where the solution is relatively large but not maximal. In this example, the time dependence is weak. The step size control tends to lead to Krylov subspaces of successively larger dimension. In this case, we need a step size control that also takes into account the Arnoldi residual. So step 5 in our algorithm becomes relevant.

The tolerance 10^{-5} might seem at first to be too small, but bear in mind that it concerns the absolute error, and the value of the solution in component 1050

is 1.5×10^{-3} . The dual solution has been computed with step size control based on an estimate similar to (10) for the dual master equation with a tolerance of 10^{-4} . We see that it does not pay off to compute the dual in this example. Looking at the error in component 1150, we see that all estimates are effective and the efficiency increases as more information from the dual is included. As expected from theory, estimate (33) also ensures that the ℓ_∞ error is below the tolerance.

In this experiment, the Moan–Niesen criterion (26) is violated in each time step. Nevertheless, we get an accurate result. This is numerical evidence that large time steps with a Magnus based integrator for the explicitly time-dependent master equation can still be very accurate. However, in this example, with $N = 2000$, the error due to the Magnus-expansion is almost insignificant compared to the Arnoldi error. Therefore, we repeated the experiment with $N = 20$. In this case, the Magnus error controls the step size and not the Arnoldi error. For each time step, we compute $\left\| \int_t^{t+\Delta t} A(\tau) d\tau \right\|_{\ell_2}$ because this is a lower bound for $\int_t^{t+\Delta t} \|A(\tau)\|_{\ell_2} d\tau$. Fig. 1 shows an estimate of the Moan–Niesen criterion for a simulation with tolerance 10^{-3} . Again, in this experiment the solution is recovered correctly (ℓ_∞ error of the size 1.3×10^{-4}) even though the Moan–Niesen criterion is violated in several steps.

Table 2: Errors and necessary number of matrix-vector products (MVPs) for $t_{\text{final}} = 10$ and tolerance 10^{-5} for the CME with varying propensity (Sec. 9.2).

estimate	ℓ_∞ error	error in comp. 1150	# MVPs (primal+dual)
(31)	$1.5 \cdot 10^{-4}$	$6.4 \cdot 10^{-6}$	28201+27191
(32)	$1.0 \cdot 10^{-5}$	$3.1 \cdot 10^{-6}$	21800+27191
(33)	$8.1 \cdot 10^{-7}$	$2.5 \cdot 10^{-7}$	31928

So far, we have computed on the whole state space. However, most of the probability mass of the solution is confined to a certain region, which is oscillating in time. We therefore repeated the experiment, but this time we allowed truncation of the state space. We start with 400 states. The adaptive control of the state space successively reduces this to about 250 states. After this initial phase, the size varies between 121 and 261 states. In this case, the problem is small enough that we can use the Matlab function `expm` to compute the matrix exponential. This (`expm`) almost conserves the ℓ_1 norm of the solution. Hence, we could just as well use the ‘probability-loss’ criterion (12) proposed by [43] to control the error due to state space truncation (cf. Sec. 3). Figure 2 compares both criteria (13) and (12). Both criteria show very similar behavior but the outflow criterion (13) is slightly overestimating in most steps. This shows that our criterion (13) is comparable to an estimate from the literature in examples where the latter can be applied.




Figure 1: Value of $\left\| \int_t^{t+\Delta t} A(\tau) d\tau \right\|_{\ell_2}$ for each time step (**), upper bound for convergence in Magnus expansion due to Moan and Niesen (—). CME example from Sec. 9.2.




Figure 2: Comparison of our (green) outflow-criterion (13) with the (blue) criterion (12) of Munsky and Khammash [43] based on the loss of probability for estimating the error due to state space truncation. CME example from Sec. 9.2 (without a Krylov approximation).

9.3 CME with constant propensities

Our adaptive algorithm also applies to examples with a constant matrix A . Table 3 summarizes results for the same example (42), except that we set $A = A_0$. The algorithm behaves similarly.

Table 3: Errors and necessary number of matrix-vector products (MVPs) for $t_{\text{final}} = 10$ and tolerance 10^{-5} for the CME with constant propensity matrix (Sec. 9.3).

estimate	ℓ_∞ error	error in comp. 1050	# MVPs (primal+dual)
(31)	$7.6 \cdot 10^{-5}$	$2.4 \cdot 10^{-5}$	1607+11364
(32)	$2.2 \cdot 10^{-6}$	$2.1 \cdot 10^{-7}$	1941+11364
(33)	$4.6 \cdot 10^{-10}$	$7.0 \cdot 10^{-11}$	2366

9.4 T Cell Homeostasis

Finally, we demonstrate our approach on a larger example, modeling T cell homeostasis [52, 37]. The master equation has time-dependent propensities

$$\begin{aligned}
\frac{d}{dt}p(n, n', t) &= \lambda_{n-1, n'}p(n-1, n', t) + \lambda'_{n, n'-1}p(n, n'-1, t) \\
&\quad + \mu_{n+1, n'}p(n+1, n', t) + \mu'_{n, n'+1}p(n, n'+1, t) \\
&\quad - (\lambda_{n, n'} + \lambda'_{n, n'} + \mu_{n, n'} + \mu'_{n, n'})p(n, n', t), \\
p(n, n', 0) &= \delta_{n10}\delta_{10n'}.
\end{aligned} \tag{46}$$

Here, $\mu_{n, n'} = n$, $\mu'_{n, n'} = n'$, $\lambda_{n, n'} = \lambda_{n, n'}(t) = 30n \left(\frac{1}{n+n'} + \frac{1}{n+1000} \right) \frac{1}{1+(\frac{t}{15})^5}$ and $\lambda'_{n, n'} = \lambda'_{n, n'}(t) = 30n \left(\frac{1}{n+n'} + \frac{1}{n'+1000} \right) \frac{1}{1+(\frac{t}{15})^5}$.

For this example, we also truncate the state space. We start with a state space that includes 0 to 29 copies of each species. The size grows quickly in the beginning, to a maximum of 59 copies of each species. Then, the size of the state space is gradually reduced with time. Figure 3(b) shows the size of the state space as a function of the time step. Figure 3(a) shows how the size of the time step changes. It is closely related to the nature of the time-dependence in the propensity matrix, which is reasonable given the expression for the error term due to the Magnus expansion. We do observe some kind of drift towards smaller time steps. This is probably due to starting from a simple Dirac delta distribution. As time evolves, the shape of the probability distribution gets more complicated, which is why the Krylov algorithm will have to do more work. A similar behaviour is also seen in the other examples above. The number of Arnoldi iterations in each time step is nearly constant (Figure 3(a)). Near the end of the simulation, the Krylov subspace is very small even though the time

steps are getting a little larger. This shows that the size of the Krylov space is not only related to the size of the time step, but also to the nature of the time-dependence.


We do not have an exact analytic solution for this example, so we use a computation with a large state space and small time steps as a ‘reference solution.’ In this way, we estimate the ℓ_∞ error of our simulation to be $2.3 \cdot 10^{-5}$. This has been achieved with 2124 time steps and 8704 matrix-vector products. If instead we take equally many time steps with a fixed step size, but allow (adaptively) variable sizes of both the state space and of the Krylov subspace, then we get an ℓ_∞ error of $7.9 \cdot 10^{-5}$, with 8289 matrix-vector products. Fixing also the size of the Krylov subspace, yields an error of $3.4 \cdot 10^{-3}$ for 8496 matrix-vector products (four per time step), or $3.5 \cdot 10^{-4}$ for 10620 matrix-vector products (five per time step). A benefit of the framework for *a posteriori* error estimation is that it makes adaptivity possible. These experiments demonstrate that allowing adaptivity in both the step size and in the Krylov subspace yields a considerably reduced error for the same amount of work.

10 Discussion


We have devised and compared three types of *a posteriori* error estimates for the master equation. Numerical tests show that the most pessimistic estimate, (10), which does not require a numerical solution of the dual problem, is the most robust. It is also the most general because it provides an error estimate for many functionals of the error at the same time, as in (11). Often, it is the most computationally efficient.

When choosing a numerical method it is desirable to mimic the conservation properties of the continuous problem. With a matrix A of the type that arises in master equations, we have ℓ_1 norm conservation, as well as positivity of the vector p . That is, $e^T p(t) \equiv 1$ and $p(t) \geq 0$ for all t . The first property is a linear invariant and is therefore preserved by any Runge–Kutta method [20, 22]. However, most Runge–Kutta methods do not respect positivity. Indeed, it has been shown that unconditionally positive Runge–Kutta methods can only be first order [7]. Uniformization is an alternative but also only of low order [49]. There have been efforts to devise other second-order schemes [9, 8]. We found this type of scheme inefficient for our examples.

The Magnus–Arnoldi propagator does not fully preserve the geometric structures of the master equation. These structures are coupled to the two properties (6) and (7) of a matrix $A(t)$ for the master equation. It would be nice if the truncated Magnus propagator shared these same properties. It is easy to verify that the first property of ‘zero column sum’ is conserved regardless the number of terms included in the approximation of the Magnus propagator. However, commutation of matrices does not preserve the second property, which is why truncated Magnus propagators of order greater than two may have negative off-diagonal entries (unlike $A(t)$). Finally, the Arnoldi process conserves the ℓ_2 but not the ℓ_1 norm of a vector. In practice, we observe that the norm-conservation



(a) Step size (red) and Krylov space (black)



(b) Size of state space

Figure 3: Choice of temporal step size, dimensionality of the Krylov subspace, and the state space for the immunology example in Sec. 9.4.

error is of considerably smaller size than the ℓ_2 - or ℓ_∞ -error of the solution.

Alternatively, one can use a numerical integrator that does not mimic the conservation properties of the continuous problem and simply project the propagated solution onto the manifold defined by the conservation properties (cf. [46]). However, this strategy is not suitable when a truncated space is used in combination with estimates such as (12) of [43] that exploit a loss of probability to estimate the error due to state space truncation.

11 Conclusions

An *a posteriori* error estimate for the master equation, based on adjoint methods, has been derived. As a general purpose estimate, it can be combined with many different types of numerical approaches to the master equation. As one example, we used the estimate to develop an adaptive algorithm that incorporates a Magnus–Arnoldi propagator into this framework. This type of evolution operator has proven to be efficient for the master equations in immunology and in chemistry that we have studied. Part of the reason for this is that relatively large time steps can be used whilst maintaining good accuracy and stability. Previous methods have employed error estimates, associated with truncation of the state space, that are valid when it is assumed that the matrix exponential is computed exactly. This is an issue because in practice the exponential is almost always approximated. The adjoint framework has allowed us to address this issue by suggesting an alternative error estimate that can be incorporated into the framework in a way that is compatible with a Krylov approximation to the matrix exponential. Overall, adjoint methods offer a mathematical framework for combining and controlling errors from different sources, giving more confidence in the final computed solution.

12 Acknowledgment

Katharina Kormann began this work in Uppsala University.

References

- [1] A. H. AL-MOHY AND N. J. HIGHAM, *Computing the action of the matrix exponential, with an application to exponential integrators*, SIAM J. SCI. COMPUT., 33 (2011), pp. 488–511, doi:10.1137/100788860.
- [2] A. ARKIN, J. ROSS, AND H. H. MCADAMS, *Stochastic kinetic analysis of developmental pathway bifurcation in phage lambda-infected Escherichia coli cells*, Genetics, 149 (1998), pp. 1633–1648, <http://www.genetics.org/content/149/4/1633>.

- [3] U. ASCHER AND L. PETZOLD, *Computer Methods for Ordinary Differential Equations and Differential-Algebraic Equations*, SIAM, 1998, <http://bookstore.siam.org/ot61/>.
- [4] S. BLANES, F. CASAS, J. A. OTEO, AND J. ROS, *The Magnus expansion and some of its applications*, Phys. Rep., 470 (2009), pp. 151–238, doi:10.1016/j.physrep.2008.11.001.
- [5] S. BLANES, F. CASAS, AND J. ROS, *Improved high order integrators based on the Magnus expansion*, BIT Numer. Math., 40 (2000), pp. 434–450, doi:10.1023/A:1022311628317.
- [6] R. A. BLYTHE AND A. J. MCKANE, *Stochastic models of evolution in genetics, ecology and linguistics*, Jour. Stat. Mech.: Theory and Experiment, (2007), pp. P07018:1–58, <http://stacks.iop.org/1742-5468/2007/i=07/a=P07018>.
- [7] C. BOLLEY AND M. CROUZEIX, *Conservation de la positivité lors de la discrétisation des problèmes d'évolution paraboliques*, RAIRO – Analyse numérique, 12 (1978), pp. 237–245, <http://eudml.org/doc/193321>.
- [8] N. BROEKHUIZEN, G. RICKARD, J. BRUGGEMAN, AND A. MEISTER, *An improved and generalized second order, unconditionally positive, mass conserving integration scheme for biochemical systems*, Appl. Num. Math., 58 (2008), pp. 319–340, doi:10.1016/j.apnum.2006.12.002.
- [9] J. BRUGGEMAN, H. BURCHARD, B. KOOI, AND B. SOMMEIJER, *A second-order, unconditionally positive, mass-conserving integration scheme for biochemical systems*, Appl. Num. Math., 57 (2007), pp. 36–58, doi:10.1016/j.apnum.2005.12.001.
- [10] T. BUI-THANH, K. WILLCOX, AND O. GHATTAS, *Model reduction for large-scale systems with high-dimensional parametric input space*, SIAM Journal on Scientific Computing, 30 (2008), pp. 3270–3288, doi:10.1137/070694855, <http://dx.doi.org/10.1137/070694855>.
- [11] Y. CAO, H. LI, AND L. PETZOLD, *Efficient formulation of the stochastic simulation algorithm for chemically reacting systems*, J. Chem. Phys., 121 (2004), pp. 4059–4067, doi:10.1063/1.1778376.
- [12] Y. CAO, S. LI, L. PETZOLD, AND R. SERBAN, *Adjoint sensitivity analysis for differential-algebraic equations: The adjoint dae system and its numerical solution*, SIAM Journal on Scientific Computing, 24 (2003), pp. 1076–1089, doi:10.1137/S1064827501380630, <http://dx.doi.org/10.1137/S1064827501380630>.
- [13] Y. CAO AND L. PETZOLD, *A posteriori error estimation and global error control for ordinary differential equations by the adjoint method*, SIAM J. Sci. Comput., 26 (2004), pp. 359–374, doi:10.1137/S1064827503420969.

- [14] D. COX AND H. MILLER, *The theory of stochastic processes*, London : Methuen, 1965, doi:10.1007/978-3-642-61943-4.
- [15] P. DEUFLHARD, W. HUISINGA, T. JAHNKE, AND M. WULKOW, *Adaptive discrete galerkin methods applied to the chemical master equation*, tech. report, Zuse Institute Berlin and Freie Universität, Berlin, 2007, doi:10.1137/070689759.
- [16] P. DEUFLHARD AND M. WULKOW, *Computational treatment of polyreaction kinetics by orthogonal polynomials of a discrete variable*, Impact of Computing in Science and Engineering, 1 (1989), pp. 269–301, doi:10.1016/0899-8248(89)90013-X.
- [17] B. DRAWERT, M. TROGDON, S. TOOR, L. PETZOLD, AND A. HELLANDER, *Molns: A cloud platform for interactive, reproducible, and scalable spatial stochastic computational experiments in systems biology using pyur-dme*, SIAM Journal on Scientific Computing, 38 (2016), pp. C179–C202, doi:10.1137/15M1014784.
- [18] J. ELF, G.-W. LI, AND X. S. XIE, *Probing transcription factor dynamics at the single-molecule level in a living cell.*, Science, 316 (2007), pp. 1191–1194, doi:10.1126/science.1141967, <http://dx.doi.org/10.1126/science.1141967>.
- [19] V. FROST AND B. MELAMED, *Traffic modeling for telecommunications networks*, Communications Magazine, IEEE, 32 (1994), pp. 70–81, doi:10.1109/35.267444.
- [20] C. W. GEAR, *Invariants and numerical methods for odes*, Physica D, 60 (1992), pp. 303–310, doi:10.1016/0167-2789(92)90246-J.
- [21] D. T. GILLESPIE, *Exact stochastic simulation of coupled chemical reactions*, J. Phys. Chem., 81 (1977), pp. 2340–2361, doi:10.1021/j100540a008.
- [22] E. HAIRER, C. LUBICH, AND G. WANNER, *Geometric numerical integration: Structure-preserving algorithms for ordinary differential equations*, Springer, 2006, doi:10.1007/3-540-30666-8.
- [23] M. HEGLAND, C. BURDEN, L. SANTOSO, S. MACNAMARA, AND H. BOOTH, *A solver for the stochastic master equation applied to gene regulatory networks*, Journal of Computational and Applied Mathematics. Special issue on evolutionary problems., 205 (2006), pp. 708–724, doi:10.1016/j.cam.2006.02.053.
- [24] A. HELLANDER AND P. LÖTSTEDT, *Hybrid method for the chemical master equation*, Journal of Computational Physics, 227 (2007), pp. 100–122, doi:10.1016/j.jcp.2007.07.020.

- [25] T. HENZINGER, M. MATEESCU, AND V. WOLF, *Sliding window abstraction for infinite markov chains*, in Computer Aided Verification, A. Bouajjani and O. Maler, eds., vol. 5643 of Lecture Notes in Computer Science, Springer Berlin / Heidelberg, 2009, pp. 337–352, http://dx.doi.org/10.1007/978-3-642-02658-4_27.
- [26] A. HILFINGER AND J. PAULSSON, *Separating intrinsic from extrinsic fluctuations in dynamic biological systems*, Proc Natl Acad Sci U S A, 108 (2011), pp. 12167–12172, doi:10.1073/pnas.1018832108.
- [27] M. HOCHBRUCK AND C. LUBICH, *On Magnus integrators for time-dependent Schrödinger equations*, SIAM J. Numer. Anal., 41 (2003), pp. 945–963, doi:10.1137/S0036142902403875.
- [28] M. HOCHBRUCK, C. LUBICH, AND H. SELHOFER, *Exponential integrators for large systems of differential equations*, SIAM J. Sci. Comput., 19 (1998), pp. 1552–1574, doi:10.1137/S1064827595295337.
- [29] A. ISERLES AND S. NØRSETT, *On the solution of linear differential equations in Lie groups*, Phil. Trans. R. Soc. Lond. A, 357 (1999), pp. 983–1019, <http://www.jstor.org/stable/55187>.
- [30] T. JAHNKE AND W. HUISINGA, *Solving the chemical master equation for monomolecular reaction systems analytically*, J. Math. Biol., 54 (2007), pp. 1–26, doi:10.1007/s00285-006-0034-x.
- [31] T. JAHNKE AND W. HUISINGA, *A dynamical low-rank approach to the chemical master equation*, Bull Math Biol, 70 (2008), pp. 2283–2302, doi:10.1007/s11538-008-9346-x.
- [32] A. K. KASSAM AND L. N. TREFETHEN, *Fourth-order time-stepping for stiff pdes*, SIAM J. Sci. Comp., 26 (2005), pp. 1214–1233, doi:10.1137/S1064827502410633.
- [33] K. KORMANN, S. HOLMGREN, AND H. KARLSSON, *Global error control of the time-propagation for the Schrödinger equation with a time-dependent Hamiltonian*, J. Comput. Sci., 2 (2011), pp. 178 – 187, doi:10.1016/j.jocs.2011.02.003.
- [34] S. C. KOU AND S. G. KOU, *Modeling growth stocks via birth-death processes*, Adv. Appl. Prob., 35 (2003), p. 641?664, <http://www.jstor.org/stable/1428212>.
- [35] T. KURTZ, *Representations of markov processes as multiparameter time changes*, Ann. Probab., 8 (1980), pp. 682–715, doi:10.1214/aop/1176994660.
- [36] M. LINDNER, *Infinite Matrices and their Finite Sections*, Birkhauser, 2006, doi:10.1007/978-3-7643-7767-0.

- [37] S. MACNAMARA AND K. BURRAGE, *Stochastic modeling of naive T-cell homeostasis for competing clonotypes via the Master equation*, Multiscale Model. Simul., 8 (2010), pp. 1325–1347, doi:10.1137/09077182X.
- [38] S. MACNAMARA, K. BURRAGE, AND R. SIDJE, *Multiscale modeling of chemical kinetics via the master equation*, SIAM Multiscale Model. & Sim., 6 (2008), pp. 1146–1168, doi:10.1137/060678154.
- [39] W. MAGNUS, *On the exponential solution of differential equations for a linear operator*, Commun. Pure Appl. Math., 7 (1954), pp. 649–673, doi:10.1002/cpa.3160070404.
- [40] P. C. MOAN AND J. NIESEN, *Convergence of the Magnus series*, Found. Comput. Math., 8 (2008), pp. 291–301, doi:10.1007/s10208-007-9010-0.
- [41] C. MOLER AND C. V. LOAN, *Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later*, SIAM Rev., 45 (2003), pp. 3–49, doi:10.1137/S00361445024180.
- [42] K. MOON, A. SZEPESSY, R. TEMPONE, AND G. ZOURARIS, *A variational principle for adaptive approximation of ordinary differential equations*, Numer. Math., 96 (2003), pp. 131–152, doi:10.1007/s00211-003-0467-8.
- [43] B. MUNSKY AND M. KHAMMASH, *The finite state projection algorithm for the solution of the chemical master equation*, J. Chem. Phys., 124 (2006), p. 044104, <http://dx.doi.org/10.1063/1.2145882>.
- [44] J. R. NORRIS, *Markov chains*, Cambridge, 1997, <http://www.statslab.cam.ac.uk/~james/Markov/>.
- [45] Y. SAAD, *Analysis of some Krylov subspace approximations to the matrix exponential operator*, SIAM J. Numer. Anal., 29 (1992), pp. 209–228, doi:10.1137/0729014.
- [46] A. SANDU, *Positive numerical integration methods for chemical kinetic systems*, J. Comput. Phys., 170 (2001), pp. 589–602, doi:10.1006/jcph.2001.6750.
- [47] T. SCHMELZER AND L. N. TREFETHEN, *Evaluating matrix functions for exponential integrators via carathéodory-fejér approximation and contour integrals*, ETNA, 29 (2007), pp. 1–18, <http://eudml.org/doc/117658>.
- [48] R. B. SIDJE, *Expokit: A software package for computing matrix exponentials.*, ACM Trans. Math. Software, 24 (1998), pp. 130–156, www.expokit.org.
- [49] R. B. SIDJE, K. BURRAGE, AND S. MACNAMARA, *Inexact uniformization method for computing transient distributions of markov chains*, SIAM Journal on Scientific Computing, 29 (2007), pp. 2562–2580, doi:10.1137/060662629, <http://dx.doi.org/10.1137/060662629>.

- [50] D. M. D. SMITH, J.-P. ONNELA, AND N. S. JONES, *Master-equation analysis of accelerating networks*, Phys. Rev. E, 79 (2009), p. 056101, doi:10.1103/PhysRevE.79.056101.
- [51] W. J. STEWART, *Introduction to the Numerical Solution of Markov Chains*, Princeton University Press, 1994, <http://press.princeton.edu/titles/5640.html>.
- [52] E. R. STIRK, C. MOLINA-PARÍS, AND H. A. VAN DEN BERG, *Stochastic niche structure and diversity maintenance in the t cell repertoire*, J. Theor. Biol., 255 (2008), pp. 237–249, <http://dx.doi.org/10.1016/j.jtbi.2008.07.017>.
- [53] G. STRANG AND S. MACNAMARA, *Functions of difference matrices are Toeplitz plus Hankel*, SIAM Review, 56 (2014), pp. 525–546, doi:10.1137/120897572.
- [54] N. VAN KAMPEN, *Stochastic processes in physics and chemistry*, Elsevier, Amsterdam, 2007, http://store.elsevier.com/Stochastic-Processes-in-Physics-and-Chemistry/N_G_-Van-Kampen/isbn-9780444529657/.
- [55] V. WOLF, R. GOEL, M. MATEESCU, AND T. A. HENZINGER, *Solving the chemical master equation using sliding windows.*, BMC Syst Biol, 4 (2010), p. 42, doi:10.1186/1752-0509-4-42.